

AWS

Kinesis

what is it?

Amazon Kinesis Data Firehose is a reliable way to stream data in near real-time. Data can be streamed to S3, Redshift or Elasticsearch.

Kinesis allows data to be streamed in real-time from a Producer to a Processor or Storage option.

This is a huge change from Batch Processing that has been the traditional way to land data from one location to another.

- **Batch Processing** - data is landed in chunks and analysed when the transfer is complete
- **Stream Processing** - streams of data pour in and don't have an end... unless you create one. This allows us to act on the data and make decisions faster

AWS

Kinesis Data Streams


what is it?

This is an example of how a streaming process can work using AWS Kinesis.

The data comes into the stream, can then be processed by AWS Kinesis Analytics or another AWS Service and then lands at its destination.

use cases

- Streaming data like website clicks and transactional data
- Migrating data from databases
- Applications with specialised data pipelines



AWS

Kinesis Firehose

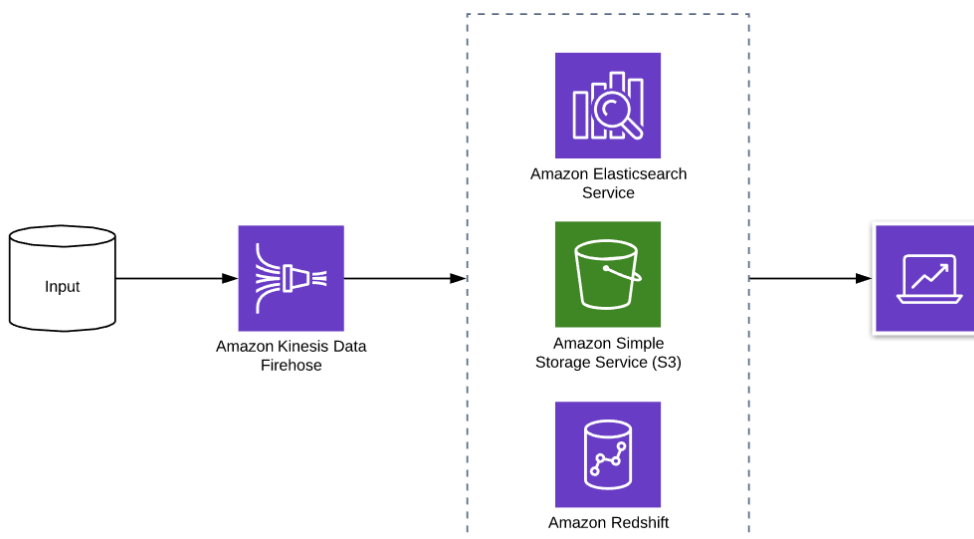
what is it?

Kinesis Firehose differs from Kinesis Data Streams as it takes the data, batches, encrypts and compresses it.

Then persists it somewhere such as Amazon S3, Amazon Redshift, or Amazon Elasticsearch.

use cases

- IoT events
- Security monitoring as Splunk can be configured as a destination
- Auto archiving



AWS

Kinesis Analytics

what is it?

Kinesis Data Analytics allows us to both process events and analyse them using SQL queries on-the-fly.

The service recognises formats like JSON and CSV, then sends the output to an analytics tool for visualisation or action.

use cases

- Processing of events data from applications
- Exploratory analysis
- Analysing clickstream anomalies

